

L16	85	(map\$4 near2 (hierarch\$4 or (multi adj level))) near3 ((logical or virtual) address\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L17	1	redirect\$4 near2 ((second\$4 adj virtual) adj (memory or volume or storage or system or network or device))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L18	250	((second\$4 adj virtual) adj (memory or volume or storage or system or network or device))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L19	250	((second\$4 adj virtual) adj (memory or volume or storage or system or network or device))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L20	20	map\$4 near2 ((second\$4 adj virtual) adj (memory or volume or storage or system or network or device))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L21	31	map\$4 near2 ((second\$4 adj virtual) near2 (memory or volume or storage or system or network or device))	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L22	2	"6735765".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:50
L23	2	"6304951".pn.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:50
L24	4	L22 or L23	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:50
L25	4	L24 or L23	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:50

L26	1	20 and 25	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:52
L27	37	1 or 2 or 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:52
L28	0	25 and 27	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:52

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	8	(Troy near2 Pesola).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:44
L2	4	(Ravi near2 Kavuri).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:45
L3	29	(Michael near2 Booth).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:45
L4	37	1 or 2 or 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:45
L5	2	1 and 2 and 3	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:45
L11	387	virtual adj3 map\$4 adj2 virtual	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L12	174	virtual adj2 map\$4 adj2 virtual	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L13	5688	(primary near2 virtual) same map\$4 (second\$3 near2 virtual)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L14	17	(primary near2 virtual) same map\$4 same(second\$3 near2 virtual)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48
L15	281	(map\$4 near2 (hierarch\$4 or (multi adj level))) same ((logical or virtual) address\$2)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/03 14:48


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction sur](#)

Terms used

map mapping or mapped virtual volume redirected redirect media type media type physical virtual logical

 Sort results by
☒ [Save results to a Binder](#)
[Try an Advanced Search](#)

 Display results
☒ [Search Tips](#)
[Try this search in The ACM Guide](#)
☐ [Open results in a new window](#)

Results 1 - 20 of 25

 Result page: [1](#) [2](#) [next](#)

 Relevance scale ☐ ☐

1 [Distributed file systems: concepts and examples](#)

Eliezer Levy, Abraham Silberschatz

 December 1990 **ACM Computing Surveys (CSUR)**, Volume 22 Issue 4

 Full text available: ☒ pdf(5.33 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

The purpose of a distributed file system (DFS) is to allow users of physically distributed computers to share data and storage resources by using a common file system. A typical configuration for a DFS is a collection of workstations and mainframes connected by a local area network (LAN). A DFS is implemented as part of the operating system of each of the connected computers. This paper establishes a viewpoint that emphasizes the dispersed structure and decentralization of both data and con ...

2 [4.2BSD and 4.3BSD as examples of the UNIX system](#)

John S. Quarterman, Abraham Silberschatz, James L. Peterson

 December 1985 **ACM Computing Surveys (CSUR)**, Volume 17 Issue 4

 Full text available: ☒ pdf(4.07 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

This paper presents an in-depth examination of the 4.2 Berkeley Software Distribution, Virtual VA: 11 Version (4.2BSD), which is a version of the UNIX Time-Sharing System. There are notes throughout on 4.3BSD, the forthcoming system from the University of California at Berkeley. We trace the historical development of the UNIX system from its conception in 1969 until today, and describe the design principles that have guided this development. We then present the internal data structures and ...

3 [A structural view of the Cedar programming environment](#)

Daniel C. Swinehart, Polle T. Zellweger, Richard J. Beach, Robert B. Hagmann

 August 1986 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 8 Issue 4

 Full text available: ☒ pdf(6.32 MB)

 Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents an overview of the Cedar programming environment, focusing on its overall structure—that is, the major components of Cedar and the way they are organized. Cedar support the development of programs written in a single programming language, also called Cedar. Its primary purpose is to increase the productivity of programmers whose activities include experimental programming and the development of prototype software systems for a high-performance personal computer. T ...

4 IP switching—ATM under IP

Peter Newman, Greg Minshall, Thomas L. Lyon

April 1998 **IEEE/ACM Transactions on Networking (TON)**, Volume 6 Issue 2

Full text available:  [pdf\(154.32 KB\)](#)


Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: Internet protocol, asynchronous transfer mode, broadband communication, communication system control, data communication, packet switching, protocols

5 System support for pervasive applications

Robert Grimm, Janet Davis, Eric Lemar, Adam Macbeth, Steven Swanson, Thomas Anderson, Brian Bershad, Gaetano Borriello, Steven Gribble, David Wetherall

November 2004 **ACM Transactions on Computer Systems (TOCS)**, Volume 22 Issue 4

Full text available:  [pdf\(1.82 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Pervasive computing provides an attractive vision for the future of computing. Computational power will be available everywhere. Mobile and stationary devices will dynamically connect and coordinate to seamlessly help people in accomplishing their tasks. For this vision to become a reality, developers must build applications that constantly adapt to a highly dynamic computing environment. To make the developers' task feasible, we present a system architecture for pervasive computing, called & ...

Keywords: Asynchronous events, checkpointing, discovery, logic/operation pattern, migration, one.world, pervasive computing, structured I/O, tuples, ubiquitous computing

6 Technical reports

SIGACT News Staff

January 1980 **ACM SIGACT News**, Volume 12 Issue 1

Full text available:  [pdf\(5.28 MB\)](#)

Additional Information: [full citation](#)

7 Conversion technology, an assessment

James P. Fry

July 1981 **ACM SIGMIS Database , ACM SIGMOD Record**, Volume 12,13 , 12 Issue 4,1 , 2

Full text available:  [pdf\(2.36 MB\)](#)

Additional Information: [full citation](#), [references](#)

8 Phoenix: a parallel programming model for accommodating dynamically joining/leaving resources

Kenjiro Taura, Kenji Kaneda, Toshio Endo, Akinori Yonezawa

June 2003 **ACM SIGPLAN Notices , Proceedings of the ninth ACM SIGPLAN symposium on Principles and practice of parallel programming**, Volume 38 Issue 10

Full text available:  [pdf\(197.86 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper proposes Phoenix, a programming model for writing parallel and distributed application that accommodate dynamically joining/leaving compute resources. In the proposed model, nodes involved in an application see a large and fixed *virtual node name space*. They communicate via messages, whose destinations are specified by virtual node names, rather than names bound to a physical resource. We describe Phoenix API and show how it allows a transparent migration of application states ...

Keywords: distributed programming, message passing, migration, parallel programming, resource reconfiguration

9 The use of connectionless network layer protocols over FDDI networks

Dave Katz

July 1990 **ACM SIGCOMM Computer Communication Review**, Volume 20 Issue 3

Full text available:  pdf(1.15 MB)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Methods for running the DoD IP and OSI connectionless network layer protocols over the FDDI medium are presented. Issues specific to the interaction between network layer protocols and FDDI are discussed, and some possible approaches to problems encountered are evaluated. The OSI protocol suite is examined in particular detail. This work was supported in part by National Science Foundation agreement no. NCR 8720904.

10 Tools and approaches for developing data-intensive Web applications: a survey

Piero Fraternali

September 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 3

Full text available:  pdf(524.80 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The exponential growth and capillar diffusion of the Web are nurturing a novel generation of applications, characterized by a direct business-to-customer relationship. The development of such applications is a hybrid between traditional IS development and Hypermedia authoring, and challenges the existing tools and approaches for software production. This paper investigates the current situation of Web development tools, both in the commercial and research fields, by identifying and characterizing ...

Keywords: HTML, Intranet, WWW, application, development

11 Innovative Document Systems: The multivalent browser: a platform for new ideas

Thomas A. Phelps, Robert Wilensky

November 2001 **Proceedings of the 2001 ACM Symposium on Document engineering**

Full text available:  pdf(188.51 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Multivalent Browser is built on a architecture that separates functionality from concrete document format. Almost all functionality is made available via relatively small modules of code called behaviors that programmers can write to extend the core system. Behaviors can be as significant and powerful as parser-renderers for scanned paper, HTML, or TeX DVI; as fine-grained as hyperlinks, cookies, and the disabling of menu items; and as innovative or uncommon as in situ annotations, "lenses", ...

Keywords: annotation, architecture, digital, document, multivalent behavior, paper, scanned

12 A cost-effective, high-bandwidth storage architecture

Garth A. Gibson, David F. Nagle, Khalil Amiri, Jeff Butler, Fay W. Chang, Howard Gobioff, Charles Hardin, Erik Riedel, David Rochberg, Jim Zelenka

October 1998 **Proceedings of the eighth international conference on Architectural support for programming languages and operating systems**, Volume 32, 33 Issue 5, 11

Full text available:  pdf(1.67 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)


This paper describes the Network-Attached Secure Disk (NASD) storage architecture, prototype implementations of NASD drives, array management for our architecture, and three filesystems built on our prototype. NASD provides scalable storage bandwidth without the cost of servers used

primarily, for transferring data from peripheral networks (e.g. SCSI) to client networks (e.g. ethernet). Increasing dataset sizes, new attachment technologies, the convergence of peripheral and interprocessor switc ...

13 Towards a universal directory service

Keith A. Lantz, Judy L. Edighoffer, Bruce L. Hitson

August 1985 **Proceedings of the fourth annual ACM symposium on Principles of distributed computing**


Full text available:  pdf(1.18 MB)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

14 Internet security attacks at the basic levels

Marco de Vivo, Gabriela O. de Vivo, Germinal Isern

April 1998 **ACM SIGOPS Operating Systems Review**, Volume 32 Issue 2

Full text available:  pdf(1.28 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

The Internet put the rest of the world at the reach of our computers. In the same way it also made our computers reachable by the rest of the world. Good news and bad news!. Over the last decade the Internet has been subject to widespread security attacks. Besides the classical terms, new one had to be found in order to designate a large collection of threats: *Worms, break-ins, hackers, crackers, hijacking, phrackers, spoofing, man-in-the-middle, password-sniffing, denial-of-service*, an ...

Keywords: Client-Server, Covert Channel, DNS, Denial of Service, Ethernet, Hijacking, ICMP, Kerberos, One-Time Password, Ping, RIP, Sniffing, Spoofing, TCP/IP

15 Towards a universal directory service

Keith A Lantz, Judy L Edighoffer, Bruce L Hitson

April 1986 **ACM SIGOPS Operating Systems Review**, Volume 20 Issue 2

Full text available:  pdf(1.01 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Directory services and name servers have been discussed and implemented for a number of distributed systems. Most have been tightly interwoven with the particular distributed systems of which they are a part: a few are more general in nature. In this paper we survey recent work in the area and discuss the advantages and disadvantages of a number of approaches. From this, we are able to extract some fundamental requirements of a naming system capable of handling a wide variety of object types in ...

16 A survey of peer-to-peer content distribution technologies

Stephanos Androutsellis-Theotokis, Diomidis Spinellis

December 2004 **ACM Computing Surveys (CSUR)**, Volume 36 Issue 4

Full text available:  pdf(517.77 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Distributed computer architectures labeled "peer-to-peer" are designed for the sharing of computer resources (content, storage, CPU cycles) by direct exchange, rather than requiring the intermediation or support of a centralized server or authority. Peer-to-peer architectures are characterized by their ability to adapt to failures and accommodate transient populations of nodes while maintaining acceptable connectivity and performance. Content distribution is an important peer-to-peer application ...


Keywords: Content distribution, DHT, DOLR, grid computing, p2p, peer-to-peer

17 Virtual machines: ReVirt: enabling intrusion analysis through virtual-machine logging and

replay

George W. Dunlap, Samuel T. King, Sukru Cinar, Murtaza A. Basrai, Peter M. Chen

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Full text available:  [pdf\(1.56 MB\)](#)


Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Current system loggers have two problems: they depend on the integrity of the operating system being logged, and they do not save sufficient information to replay and analyze attacks that include any non-deterministic events. ReVirt removes the dependency on the target operating system by moving it into a virtual machine and logging below the virtual machine. This allows ReVirt to replay the system's execution before, during, and after an intruder compromises the system, even if the intruder rep ...

18 The Outlaw 'Net': Opposition to ICANN's New Internet Order

Enda Brophy

December 2002 **ACM SIGCAS Computers and Society**, Volume 32 Issue 4


Full text available:  [htm\(132.34 KB\)](#)

Additional Information: [full citation](#), [index terms](#)

19 Risks to the public: Risks to the public in computers and related systems

Peter G. Neumann

May 2004 **ACM SIGSOFT Software Engineering Notes**, Volume 29 Issue 3

Full text available:  [pdf\(128.46 KB\)](#) Additional Information: [full citation](#)

20 The RADIANCE lighting simulation and rendering system

Gregory J. Ward

July 1994 **Proceedings of the 21st annual conference on Computer graphics and interactive techniques**

Full text available:  [pdf\(2.36 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper describes a physically-based rendering system tailored to the demands of lighting design and architecture. The simulation uses a light-backwards ray-tracing method with extensions to efficiently solve the rendering equation under most conditions. This includes specular, diffuse and directional-diffuse reflection and transmission in any combination to any level in any environment, including complicated, curved geometries. The simulation blends deterministic and stochastic ray-tracing ...

Keywords: Monte Carlo, lighting simulation, physically-based rendering, radiosity, ray-tracing

Results 1 - 20 of 25

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)